

PATENT APPLICATION

REDIRECTION TO A VIRTUAL ROUTER

Inventors: Ian Wilson
1 Westhall Gardens
Edinburgh
EH10 4JJ
U.K.

Citizenship: British

Assignee: Cisco Technology, Inc.
170 West Tasman Drive
San Jose, California 95134-1706

A corporation of California

Prepared by:

BEYER, WEAVER & THOMAS, LLP

P.O. Box 778

Berkeley, CA 94704-0778

REDIRECTION TO A VIRTUAL ROUTER

BACKGROUND OF THE INVENTION

1. Field of the Invention

5 The present invention relates to networking technology. More particularly, the present invention relates to providing ICMP redirects in a router implementing a virtual router protocol.

2. Description of the Related Art

10 Networks are commonly used by organizations for a variety of purposes. For instance, through the use of networks, resources such as programs and data may be shared by users of the network. In addition, a computer network can serve as a powerful communication medium among widely separated users.

15 Communication among hosts and users of a network is often facilitated through connection to one or more routers. As shown in FIG. 1, a host 112 normally resides on a network segment 114 that enables its network entities to communicate with other entities or networks. Note that the host 112 need not directly connect to the entities or networks with which it communicates. For example, as shown in FIG. 1, the host 112 may be connected through one of several routers R1 116, R2 118, or R3 120.

20 Now, suppose that the host 112 wishes to send a message to a corresponding node

122. A message from the host 112 is then packetized and forwarded through router R1 116 and to the corresponding node 122 according to a standard protocol. If the corresponding node 122 wishes to send a message to the host 112 – whether in reply to a message from the host 112 or for any other reason – it addresses that message to the IP address of the host 112 on the network segment 114. The packets of that message are then forwarded to router R1 116 and ultimately to the host 112.

As described above, packets sent to and from the corresponding node 122 are forwarded via the router R1 116. As shown, the router R1 116 is the only route to and from the corresponding node 122. Thus, if the router R1 116 fails, communication with the corresponding node 122 becomes impossible. Accordingly, the reliability of the network as well as the routers in the network is of utmost importance.

As networks become a critical resource in many organizations, it is important that the networks are reliable. One way of achieving reliability is through redundancy. As described above, a single router failure may prevent communication to and from each host and user connected to the router. In many networks, it is common to provide redundancy through the use of multiple routers such that a backup router functions in the event of failure of a primary router. This is accomplished through the use of a virtual router protocol such as the Virtual Router Redundancy Protocol (VRRP) or the Hot Standby Redundancy Protocol (HSRP). HSRP is further described in U.S. Patent No. 5,473,599.

According to HSRP, a protocol available from Cisco Systems, Inc. located in San Jose, California, multiple routers share a common MAC address and a virtual IP address. In addition, each of the routers has its own unique MAC address that will be used by the router for advertising routes to other routers. A router may be in one of three states

according to HSRP: an active state, a standby state, and a listen state. When a router is in the active state, the router is responsible for actively forwarding packets addressed to the virtual IP address. Similarly, when a router is in the standby state, the router becomes active when the active router becomes non-functional. Thus, when the active router fails, the standby router takes over both the shared virtual IP address and the shared MAC address. A router that is in the listen state merely listens to packets that are sent to it and collects this information in the event that it later becomes an active or standby router.

As shown in FIG. 2, a system diagram illustrates the state of each router in accordance with the HSRP protocol. In this example, a network segment "Net C" 124 is connected to the second router R2 118 and the third router R3 120. When the host 112 wishes to send a packet to the subnet "Net C" 124, it may send the packet via the second router R2 118 or the third router R3 120. The Internet Control Message Protocol (ICMP), described in RFC 792 of the Network Working Group, J. Postel, September 1981, may be used to notify the host that the optimal path to the subnet 124 is via the second router R2 118. More particularly, when a gateway such as router R1 116 receives a packet, it checks its routing table to obtain the address of the next hop, or next gateway, router R2 118. If the second router R2 118 and the host 112 are on the same network, a redirect message is sent to the host 112. The redirect message advises the host 112 to send its traffic for Network C directly to the second router R2 118 since this is a shorter path to the destination. This is accomplished by specifying the IP address of the second router R2 118 in the redirect message.

While the use of an ICMP redirect message may be used to notify a router of the optimal path to a particular destination, the ICMP redirect feature is typically disabled

when a virtual router protocol such as HSRP is enabled. For example, as described above with reference to FIG. 2, a redirect message may be sent to the host 112 indicating that the optimal path to the subnet C is via the second router R2 118. Once the host 112 has discovered the IP address of the second router R2 118, it will automatically address each packet sent to the subnet C to the second router R2 118. It follows that when the second router R2 118 fails, packets sent to the second router R2 118 will never be received by the intended destination. Although the HSRP protocol is designed to send the packet via the alternate route to the subnet C 124 through the third router R3 120 upon failure of the second router R2 118, the ICMP redirect feature prevents the redundancy feature from functioning properly. Thus, although the redundancy feature is designed to ensure that the packet is received by its intended destination, the packet will not necessarily be received by its intended destination when the ICMP redirect feature is enabled. For this reason, the ICMP redirect feature is typically disabled when the HSRP protocol is enabled.

While disabling the ICMP redirect feature may be desirable in certain instances, disabling the ICMP redirect feature is problematic on networks where different routers (and/or HSRP groups) provide different routes to remote networks. More particularly, packets sent to devices on remote networks for which the HSRP active router does not provide the optimal route will traverse the local network twice since the active router must forward them to the optimal path router.

In view of the above, it would be desirable to enable a redirect feature such as the ICMP redirect feature to be used with a virtual router protocol such as HSRP. In this manner, a packet may be sent via an optimal path while retaining the redundancy features of a virtual router protocol.

SUMMARY OF THE INVENTION

The present invention enables a redirect packet to be sent in a system that
5 implements a virtual router protocol. This is accomplished, in part, by providing a virtual
IP address in the redirect packet rather than a primary IP address. In this manner, a host
may be notified that packets are to be redirected to the virtual IP address.

In accordance with one aspect of the invention, the present invention provides
methods and apparatus for sending a redirect packet to a host from a first router that
supports a virtual router protocol. The redirect packet notifies the host that specific
packets are to be redirected to a second router. One method includes receiving a packet
from a host, where the packet includes a source address identifying the host and a
destination address identifying a destination network. The first router ascertains the
destination network of the packet from the destination address and obtains from a routing
table an address of a next router to the packet's destination network. The first router then
determines whether to send a redirect packet to the host. In accordance with one
embodiment, this is performed by determining whether the next router and the host
identified by the source address of the packet are on the same network. When it is
determined that the next router and the host are on the same network, the first router
composes and sends the redirect packet to the host. The redirect packet serves to notify
the host that packets addressed to the destination network are to be redirected to a virtual
address of the next router.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a diagram illustrating communication between a host and a corresponding
5 node via a router.

FIG. 2 is a diagram illustrating a system in which multiple routers are used to
provide redundancy.

FIG. 3 is a diagram illustrating an exemplary redirect packet that provides a virtual
IP address of a router supporting a virtual router protocol in accordance with an
10 embodiment of the invention.

FIG. 4 is a diagram illustrating an exemplary virtual router table that identifies one
or more virtual IP addresses for one or more routers in accordance with one embodiment
of the invention.

FIG. 5 is a process flow diagram illustrating a method of sending a redirect packet
15 in a router that supports a virtual router protocol.

FIG. 6 is a block diagram of a network device that may be configured to
implement aspects of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

In the following description, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art, that the present invention may be practiced without some or all of these specific details. In other instances, well known process steps have not been described in detail in order not to unnecessarily obscure the present invention.

As described above, when a redirect message is received by a router that supports a virtual router protocol, the redirect message may defeat the purpose of that the virtual router protocol. More particularly, when a redirect packet (e.g., ICMP redirect) is typically sent, it provides an IP address of a router. In other words, when a virtual router protocol such as HSRP is running on the router, a primary IP address of the router rather than the virtual IP address is typically provided. As a result, the redirect feature is typically disabled when a virtual router protocol such as HSRP is in use.

In accordance with the present invention, an improved redirect message provides a virtual IP address of the router to which packets are to be redirected (e.g., an active router for the virtual router protocol). FIG. 3 is a diagram illustrating an exemplary redirect packet that provides a virtual IP address of a router supporting a virtual router protocol in accordance with an embodiment of the invention. As shown, a redirect packet 302 identifies a destination network 304 (e.g., network C) and a next hop router 306 that is identified by a virtual IP address. When the redirect packet is received by the host, the host is notified that packets addressed to the destination network 304 are to be redirected to the virtual IP address 306.

A router may send a redirect packet by accessing a virtual router table that associates one or more virtual IP addresses with each destination router to which packets

may be redirected. FIG. 4 is a diagram illustrating an exemplary virtual router table that identifies one or more virtual IP addresses for one or more routers in accordance with one embodiment of the invention. As shown in FIG. 4, a virtual router table 402 may be maintained by each router so that a virtual IP address associated with a particular destination router may be obtained and provided in a redirect packet. As shown, a virtual router table may identify an IP address 404 for a router. For instance, the IP address may be a primary IP address of an active router. One or more virtual IP addresses 406 are then associated with the IP address 404. In this example, a first router R1 is associated with two virtual IP addresses, while a second router R2 is associated with a single virtual IP address. Similarly, a third router R3 is shown to be associated with three virtual IP addresses. Thus, a router sending a redirect packet to a host may notify the host to redirect packets to any one of these virtual IP addresses.

FIG. 5 is a process flow diagram illustrating a method of sending a redirect packet in a router that supports a virtual router protocol. A first router receives a packet from a host and addressed to a remote subnet at block 502. Since the packet contains a header including a source address identifying the host and a destination address identifying a destination host, the router is able to ascertain the destination network of the packet from the destination address. The router then determines at block 504 whether a redirect message must be sent. More particularly, this is accomplished by obtaining an address of the next hop to the packet's destination network from a routing table. The router then determines whether this next hop router and the host identified by the source address of the packet are on the same network. In other words, the router determines from the routing table that the next hop to reach the remote subnet is back out the same interface

that the packet was received on.

When it is determined that the next hop router and the host are on the same network, a redirect packet is composed and sent to notify the host that packets addressed to the destination network are to be redirected to a primary or virtual address of the next hop router. In this manner, the router composes a redirect message such as an ICMP redirect packet that contains another gateway that will provide a more optimal path for the host to the remote subnet. The composition and sending of a redirect message that contains another gateway that will provide a more optimal path to the remote subnet is described in further detail below with reference to blocks 506 through 524.

In accordance with one embodiment, the router determines whether to send a redirect packet to a primary IP address of the next hop router or to a virtual IP address of the next hop router. More particularly, as shown at block 506, the router determines whether a virtual router protocol is running on the router responsible for composing and sending the redirect message. For instance, the router responsible for composing and sending the redirect message may determine whether a virtual router protocol is running on the interface on which the packet was received. If a virtual router protocol is not running on the interface of the router on which the packet was received, a redirect message including a primary IP address of the next hop router is composed and sent to the host as shown at block 508. Otherwise, the router determines at block 510 whether one or more HSRP groups are configured for the next hop router. When it is determined that no HSRP groups are configured for the next hop router, the redirect packet is composed and sent to the primary IP address of the next hop router as described above with reference to block 508. Otherwise, the router performs a look up in a virtual router table at block 512 for a

primary IP address of the next hop router.

When the primary IP address of the next hop router is found in the virtual router table, one or more virtual IP addresses may then be obtained. More particularly, the router determines at block 514 whether the next hop router is an active router for the virtual router protocol. When it is determined that the next hop router is an active router for the virtual router protocol, a redirect packet containing a virtual address of the next hop router is composed and sent to the host. More particularly, as described above with reference to FIG. 4, more than one HSRP group of a router may simultaneously be in an active state. Thus, as shown at block 516, when the next hop router has more than one group in an active state, a redirect message including one of the virtual IP addresses is composed and sent to the host. For example, the virtual IP address of the first group in an active state that is found in the virtual router table may be provided in the redirect message. However, when the gateway is not an active router for the virtual router protocol (i.e., the router has no HSRP groups in the active state), a redirect message is not sent to the host as shown at block 522. Thus, a redirect will not be sent if it would result in a host learning a router's primary address.

A software or software/hardware hybrid redirection system is preferably implemented on a general-purpose programmable machine selectively activated or reconfigured by a computer program stored in memory. Such programmable machine may be a network device designed to handle network traffic. Such network devices typically have multiple network interfaces including frame relay and ISDN interfaces, for example. Specific examples of such network devices include routers and switches. For example, the redirection systems of this invention may be specially configured routers

such as specially configured router models 1600, 2500, 2600, 3600, 4500, 4700, 7200, 7500, and 12000 and Catalyst switches such as models 5000 and 6000 available from Cisco Systems, Inc. of San Jose, California. A general architecture for some of these machines will appear from the description given below. In an alternative embodiment, the redirection system may be implemented on a general-purpose network host machine such as a personal computer or workstation. Further, the invention may be at least partially implemented on a card (e.g., an interface card) for a network device or a general-purpose computing device.

Referring now to FIG. 6, a router 1440 suitable for implementing the present invention includes a master central processing unit (CPU) 1462, interfaces 1468, and a bus 1415 (e.g., a PCI bus). When acting under the control of appropriate software or firmware, the CPU 1462 is responsible for such router tasks as routing table computations and network management. It may also be responsible for functions previously described, such as virtual router table databases, modifying tables of the redirecting router, etc. It preferably accomplishes all these functions under the control of software including an operating system (e.g., the Internetwork Operating System (IOS®) of Cisco Systems, Inc.) and any appropriate applications software. CPU 1462 may include one or more processors 1463 such as a processor from the Motorola family of microprocessors or the MIPS family of microprocessors. In an alternative embodiment, processor 1463 is specially designed hardware for controlling the operations of router 1440. In a specific embodiment, a memory 1461 (such as non-volatile RAM and/or ROM) also forms part of CPU 1462. However, there are many different ways in which memory could be coupled to the system.

The interfaces 1468 are typically provided as interface cards (sometimes referred to as "line cards"). Generally, they control the sending and receiving of data packets over

the network and sometimes support other peripherals used with the router 1440. Among the interfaces that may be provided are Ethernet interfaces, frame relay interfaces, cable interfaces, DSL interfaces, token ring interfaces, and the like. In addition, various very high-speed interfaces may be provided such as fast Ethernet interfaces, Gigabit Ethernet interfaces, ATM interfaces, HSSI interfaces, POS interfaces, FDDI interfaces and the like. Generally, these interfaces may include ports appropriate for communication with the appropriate media. In some cases, they may also include an independent processor and, in some instances, volatile RAM. The independent processors may control such communications intensive tasks as packet switching, media control and management. By providing separate processors for the communications intensive tasks, these interfaces allow the master microprocessor 1462 to efficiently perform routing computations, network diagnostics, security functions, etc.

Although the system shown in FIG. 6 is one specific router of the present invention, it is by no means the only router architecture on which the present invention can be implemented. For example, an architecture having a single processor that handles communications as well as routing computations, etc. is often used. Further, other types of interfaces and media could also be used with the router.

Regardless of network device's configuration, it may employ one or more memories or memory modules (including memory 1461) configured to store program instructions for the general-purpose network operations and other redirection functions described herein. The program instructions may control the operation of an operating system and/or one or more applications, for example. The memory or memories may also be configured to store routing tables, virtual router table databases, etc.

Because such information and program instructions may be employed to

implement the systems/methods described herein, the present invention relates to machine readable media that include program instructions, state information, etc. for performing various operations described herein. Examples of machine-readable media include, but are not limited to, magnetic media such as hard disks, floppy disks, and magnetic tape; optical
5 media such as CD-ROM disks; magneto-optical media such as floptical disks; and hardware devices that are specially configured to store and perform program instructions, such as read-only memory devices (ROM) and random access memory (RAM). The invention may also be embodied in a carrier wave travelling over an appropriate medium such as airwaves, optical lines, electric lines, etc. Examples of program instructions include both machine code, such as produced by a compiler, and files containing higher
10 level code that may be executed by the computer using an interpreter.

Although illustrative embodiments and applications of this invention are shown and described herein, many variations and modifications are possible which remain within the concept, scope, and spirit of the invention, and these variations would become clear to
15 those of ordinary skill in the art after perusal of this application. For instance, although the specification has described routers, other entities used to redirect packets to a virtual router on remote network segments can be used as well. For example, bridges or other less intelligent packet switches may also employ the standby protocol of this invention. Accordingly, the present embodiments are to be considered as illustrative and not
20 restrictive, and the invention is not to be limited to the details given herein, but may be modified within the scope and equivalents of the appended claims.